DETECTING POLYGONS OF VARIABLE DIMENSION IN OVERHEAD IMAGES WITH PARTICLE FILTERS

Siddharth Manay and David W. Paglieroni
{manay2, paglieroni1}@llnl.gov
Lawrence Livermore National Laboratory[1]

ABSTRACT

We extend the particle filtering approach to detecting polygonal structures of fixed shape and variable size in overhead images to the more practically useful case of unconstrained side lengths. Our approach deals with multiple candidate sides by validating and clustering particles based on evidence that corners of proper acuteness and orientation (as constrained by the polygon model) might exist. A queue-regulated tracking algorithm that handles multiple candidates across multiple sides is discussed. Compelling detection results in overhead images that involve entire families of polygonal structure are provided.

Index Terms— polygon / building detection, particle filter

1. INTRODUCTION

The ability to efficiently detect polygon structures in images is of great importance for understanding overhead images containing man-made objects (buildings, vehicles, parking lots, fields…). A particle filtering approach to detecting fixed-shape polygon structures of variable position, size, and orientation in images was presented in [7]. We extend the particle filtering approach to polygon structures for which side lengths are unconstrained. This powerful new extension allows us to address the ambitious and more practically useful goal of quickly and automatically detecting entire families of polygons in overhead scenes – e.g., all L-shaped buildings with legs of any length or thickness (rather than just all L-shaped buildings with fixed relative side lengths as in [7]).

Objects (polygonal or not) of fixed size and shape can be detected in overhead images by projecting edges of 2D or 3D physical models onto the image and matching gradient directions along projected edges to pixel gradient directions [4,5]. For polygonal objects of fixed shape but variable size, matching is normally attempted by assembling collections of primitive features such as edges / lines as in [2,3], or corners / geometric invariants as in [1,9,10]. As opposed to other methods in which polygon are detected by making a single attempt to match models to images, the efficient particle filtering approach in [7] combines many attempts, based on sequential Monte Carlo sampling, to track a single polygon boundary.

Our method extends [7] to cases in which polygon side lengths are unknown. In this case, particle filtering produces multi-modal distributions of particles in which each mode represents a likely side. Weighted ensemble averages of particle tracks (the "blended" tracks in [7]) cannot be used to estimate true polygon boundaries when the distributions of particles are multi-modal - and there is no assurance that each mode will be well-sampled. Our contribution is to address these issues by using the corner similarity measure in [6] (which also contributes to particle importance weights) to form particle clusters *during* particle filtering. This clustering method naturally determines the correct number of clusters and does not depend on initial conditions or parameters.

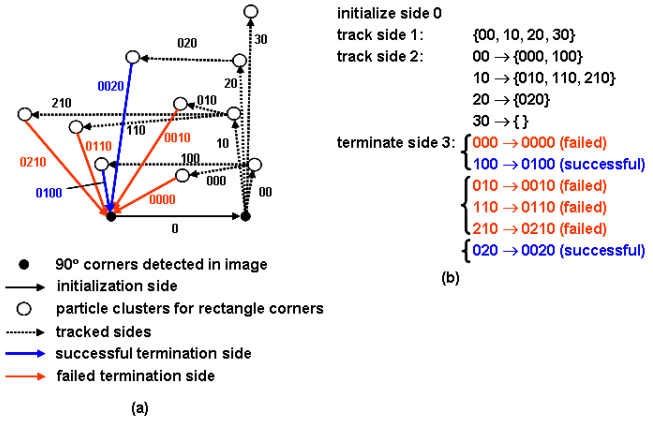

Fig.1 (a) Boundary tracking example for the model of a rectangle with unconstrained side lengths. (b) Boundary tracking detail.

Fig.1 schematically illustrates the three stages of boundary tracking (initialization, tracking, and termination) for polygons with unconstrained side lengths on a rectangle. As described in Section 2, side 0 (the solid black arrow labeled "0") is initialized from a pair of corners detected in the image (the solid dots) that satisfy geometry constraints on acuteness / orientation imposed by the rectangle model.

For iteration $0 < k < n-1$, side k (depicted as a dashed black arrow terminating on a hollow dot) is tracked by

[1] Prepared by LLNL under Contract DE-AC52-07NA27344. LLNL-CONF-400576

seeking N samples of a distribution of possible sides constrained by the model (see Section 3). Hollow dots represent clusters of closely spaced termination points (see Section 3.1) for arrows emanating either from the initialization side (the solid dot) or a particle cluster (a hollow dot).

Each arrow in Fig.1 is labeled with a sequence of from 1 to n digits. Digit k from the right identifies a specific candidate for side k. All sides except instance 3 of side 1 (labeled "30") belong to tracks that complete the polygon (the track whose last leg is labeled "30" dies). As discussed in Section 3.3, the last side (formed by connecting corner $n-1$ to corner 0 and depicted as a solid red or blue arrow) is accepted only if it nearly satisfies the geometry constraint (in which case, the arrow is blue). A queue-regulated tracking algorithm that handles multiple particle clusters across multiple sides is discussed in Section 3.2, and several detection examples are provided in Section 4.

2. POLYGON BOUNDARY TRACK INITIALIZATION

Fields of pixel gradient directions $\{\theta(c,r)\}$ at pixels with column and row coordinates $[c,r]$ can be estimated from pixel gray values $\{u(c,r)\}$ using Sobel operators (as in our work) or various other methods (e.g., Canny's method). Corners are detected at locations of strong match between pixel gradient directions and gradient directions normal to edges in models of corners and sides [6].

Following [1], we model n-sided polygons as sequences $\mathcal{P}^* = \{\alpha_k^* \quad k = 0\ldots n-1\}$, counterclockwise along the boundary, of n corners of acuteness $\alpha_k^* \in (0,2\pi) \neq \pi$ ($< \pi$ for concave corners). For $k = 0\ldots n-1$, corner k is defined by its vertex location $[c_k, r_k]$, acuteness α_k, and bisector pointing direction $\theta_k \in [0,2\pi)$. We shall define side k in terms of the location / acuteness $[c_k, r_k, \alpha_k]$ of corner k and the pointing direction / length $[\phi_k, L_k]$ of the ray that emanates from corner k and terminates on corner $(k+1) \bmod n$. The bisector pointing direction θ_k for corner k relates to the corner k acuteness α_k and the side k pointing direction ϕ_k as

$$(1)\quad \theta_k = (\phi_k - i\alpha_k/2) \bmod 2\pi\,,\ i = \begin{cases} -1 & \alpha_k < \pi \\ +1 & \alpha_k > \pi \end{cases}$$

The state vector $\boldsymbol{x_k} \triangleq [c_k, r_k, \alpha_k, \phi_k, L_k]^{\mathrm{T}}$ characterizes side k geometrically. An initial side (side 0 in Fig.1) is defined by a pair of detected corners that satisfy geometry constraints on acuteness / orientation imposed by the polygon model, and side length:

$$(2)\quad \alpha_j = \min(\alpha_j^*, 2\pi - \alpha_j^*),\quad j = 0,1$$

$$(3)\quad \theta_0 + \alpha_0^*/2 \approx \theta_1 - \alpha_1^*/2$$

$$(4)\quad L_{\min} \le L_0 \le L_{\max}$$

$L_{\min}$ and $L_{\max}$ are broad constraints on minimum and maximum allowable side lengths. A more thorough initialization is achieved by applying (2)-(4) to n cyclic shifts of polygon model $\mathcal{P}^*$, in which case, boundaries of polygons for which *any* side is represented in the field of detected corners will be tracked.

3. BOUNDARY TRACKS FOR POLYGONS WITH UNCONSTRAINED SIDE LENGTHS

Boundary trackers that detect polygons with unconstrained side lengths are described below. We introduce a novel clustering technique based on photometric evidence of corners corresponding to unknown side lengths, and then describe how boundary tracks are terminated.

3.1. Clusters of strong corner similarities

In the particle filtering approach, pointing directions ϕ_k and side lengths L_k for random samples of rays along side k emanating from corner $[c_k, r_k, \alpha_k]$ are obtained by sampling *proposed state transition densities* of ϕ and L values (Fig.2a):

$$(5)\quad \phi_k \sim q(\phi \mid \phi_{k-1}) \text{ with mode } \phi_k^* = \phi_{k-1} + \pi - \alpha_k^*$$

$$(6)\quad L_k \sim q(L) = \mathrm{U}(L_{\min}, L_{\max})$$

When side lengths are not constrained by the model, the resulting distribution of weighted (see (10)) particles $\boldsymbol{x_k^i}$, $i = 1\ldots N$ is likely to be multi-modal, with peaks at multiple likely side termination points. However, (a) the modes may be poorly sampled, (b) many particles will be "wasted" because they are far from modes, and (c) ensemble averages do not correspond to modes. We shall thus use photometric information to cluster samples in advance of particle filtering.

A similarity measure $s(c, r, \theta \mid \alpha) \in [0,1]$ sensitive to corners of acuteness α and orientation θ at vertex pixel location $[c,r]$ was proposed in [6]. For each corner acuteness α in polygon model $\mathcal{P}^*$, we search the 3D array $\{s(c, r, \theta \mid \alpha)\}$ of corner similarities for clusters of at least m 8-connected points $[c,r,\theta]$ (say m>10) that satisfy

$$(7)\quad [c,r,\theta]:\ s(c,r,\theta \mid \alpha) \ge s_{\min}$$

As discussed in [6], there are N_θ quantized corner orientations θ. $s_{\min}$ is either user specified or the minimum similarity over all detected corners. Only points $[c,r,\theta \mid \alpha]$ that satisfy (7) are assigned cluster labels. Values of cluster labels are stored in 3D arrays $\{Label(c,r,\theta \mid \alpha)\}$. The labels are pre-computed once.

N attempts are made to track each side k candidate by sampling the state transition densities in (5)-(6) N times:

$$(8)\quad \phi_k^i \sim q(\phi \mid \phi_{k-1}^i), \quad L_k^i \sim q(L), \quad i = 1...N$$

Attempt i for a particular side k candidate is the ray of length L_k^i and pointing direction ϕ_k^i emanating from $[c_k^i, r_k^i]$. The associated particle is $\boldsymbol{x_k^i} = [c_k^i, r_k^i, \alpha_k^i, \phi_k^i, L_k^i]^{\mathrm{T}}$. For each particle, the predicted corner $[c_{k+1}^i, r_{k+1}^i, \theta_{k+1}^i, \alpha_{k+1}^i]$ at the endpoint of side k is computed from $\boldsymbol{x_k^i}$ and $\mathcal{P}^*$. If $Label(c_{k+1}^i, r_{k+1}^i, \theta_{k+1}^i \mid \alpha_{k+1}^*)$ is unassigned, the particle is rejected and another particle is drawn using (8). Otherwise the particle is labeled with the assigned value. All particles with the same label are said to be part of the same cluster. We impose a limit on the number of times a particle can be re-drawn. More efficient sampling schemes are under consideration.

When side lengths are constrained, the approach in [7] copes with missing sides and corners by incorporating both geometric information (from the model) and photometric information (from image pixel gray values). However, when side lengths are unconstrained, the approach developed in this paper is more sensitive to missing information because in this more challenging case, boundary tracking must necessarily stop when photometric evidence of corners or sides is missing. In this case, no clusters are found and tracking stops.

Each particle $\boldsymbol{x_k^i}$ has an associated *importance weight* w_k^i which varies from 0 to 1. The weights are normalized for any side k so as to sum to one within a cluster. The popular *sample-importance-resample* (*SIR*) particle filter [8] computes particle weights as

$$(10)\quad w_k^i = q(\boldsymbol{z_k} \mid \boldsymbol{x_k^i}) = q(\boldsymbol{z}(\boldsymbol{x_k^i}))$$

where $\boldsymbol{z} = \boldsymbol{z}(\boldsymbol{x})$ is the feature vector of measurements (observations) made from the image, and $q(\boldsymbol{z} \mid \boldsymbol{x})$ is a *proposed measurements density*. Following [7], we use the single feature

$$(11)\quad \boldsymbol{z} = \boldsymbol{z}(\boldsymbol{x}) = s(\boldsymbol{x}) = s(c, r, \theta \mid \alpha) \in [0,1]$$

and the associated measurements density in Fig.2b, where θ is a function of α and ϕ (see (1)). The SIR filter resamples particles on each side to prevent particle weights from degenerating into a single dominant value as tracking progresses. For each cluster at each iteration $k = 1...n-2$, N resampled particles $\{\tilde{\boldsymbol{x}}_{\boldsymbol{k}}^{\boldsymbol{i}}\}_{i=1}^N$ are drawn from the subset of weighted particles $\{\boldsymbol{x_k^i}\}_{i=1}^N$, $\{w_k^i\}_{i=1}^N$ belonging to that cluster. Particles with higher weights are often selected multiple times, and particles with lower weights are often not selected at all (and are thus eliminated). Importance weights $\{\tilde{w}_k^i\}_{i=1}^N$ associated with resampled particles are inherited from the original particles and subsequently normalized.

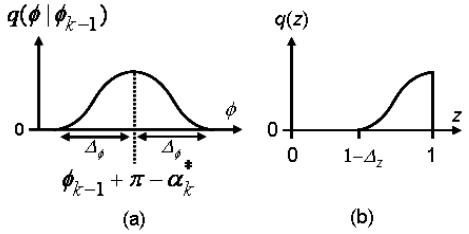

Fig.2 Proposed state transition and measurements densities.

3.2. Queue-regulated multi-cluster tracking

Each particle cluster is resampled independently of all other clusters, and in this sense is a separate particle filter. For ease of implementation, clusters can be maintained in a queue. The initialization side is en-queued. Then for each iteration $k = 1...n-2$ a cluster is de-queued, one particle filtering iteration is applied, the particles are clustered, and each cluster is resampled and en-queued. When iteration $k = n-2$ is complete, the queue will contain clustered sets of partial tracks missing only the last side.

3.3. Polygon boundary track termination

From (5), the criterion for successful termination of a polygon boundary track can be expressed as

$$(12)\quad \phi_{n-1} \approx \phi_{n-2} + \pi - \alpha_{n-1}^*$$

i.e., the actual direction of the ray pointing from corner $[c_{n-1}, r_{n-1}]$ to corner $[c_0, r_0]$ cannot deviate much from the direction predicted by the model. Once successfully tracked, the boundary tracks in each cluster are blended into a single "expected" boundary. This process can incorporate not just corner similarity, but also edge similarity measurements. The resulting boundaries are ranked with a gradient direction measure ([7]).

Fig.3 Rectangles (green) and L-shaped polygons (red) with unconstrained sides on an industrial area overhead image.

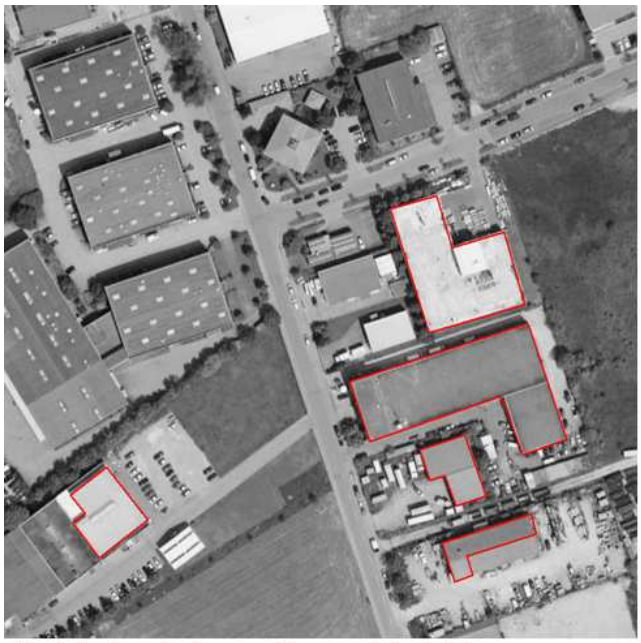

Fig.4 L-shaped polygons with unconstrained sides on an industrial area overhead image.

4. EXAMPLES

Fig.3-4 demonstrate the unconstrained side-length polygon boundary tracking algorithm for building detection in overhead images. Detected polygons are filtered using a region homogeneousness measure and disambiguated. Top ranked polygons are displayed. Disambiguation deletes polygons that are too close (based on a distance tolerance) to a higher ranked polygon; distance is based on a Hausdorff measure. Fig.3 contains the top 43 rectangles (in green) and top 5 "L"-shaped polygons (in red) detected using $N = 50$ in a 1024x1024 image of an industrial park (courtesy of CaSIL). Shapes with multiple zero-contrast edges/corners are not detected. Fig.4 contains the top 5 "L"-shaped polygons detected using $N = 100$ in a 768x768 image of an industrial park (courtesy of Google Earth.)

6. SUMMARY AND CONCLUSIONS

We have presented a novel polygon detection algorithm that uses a sequential Monte Carlo approach to track polygon boundaries without imposing constraints on the side lengths. The lack of constraints results in multi-modal particle distributions representing probable side lengths. To efficiently sample and detect these modes, particles are clustered based on the connected components of the underlying corner-similarity array. We outlined the clustering method and a queue-regulated boundary tracking algorithm, and then presented detection results on real overhead images.

5. REFERENCES

[1] S. Manay and D. Paglieroni, "Matching Flexible Polygons to Fields of Corners Extracted from Images", Int. Conf. Image Anal. and Recog. (ICIAR), August 22-24, 2007, Montreal, Canada.

[2] J. McGlone and J. Shufelt, "Projective and Object Space Geometry for Monocular Building Extraction", CVPR 1994, pp.54-61.

[3] S. Noronha and R. Nevatia, "Detection and Modeling of Buildings from Multiple Aerial Images", IEEE Trans. PAMI, vol.23, no.5, May 2001, pp.501-518.

[4] C. Olson and D. Huttenlocher, "Automatic Target Recognition by Matching Oriented Edge Pixels", IEEE Trans. Image Processing, vol.6, no. 1, January 1997, pp.103-113.

[5] D. Paglieroni, W. Eppler and D. Poland, "Phase Sensitive Cueing for 3D Objects in Overhead Images", SPIE Defense and Security Symposium: Signal Processing, Sensor Fusion and Target Recognition XIV, Proc. SPIE, Vol.5809, 28-30 March 2005, Orlando FL.

[6] D. Paglieroni and S. Manay, "Signal-to-Noise Behavior for Matches to Gradient Direction Models of Corners in Images", SPIE Defense & Security Symposium, 9-13 April 2007, Orlando, FL, ATR XVII, Advanced Algorithms for ATR.

[7] D. Paglieroni and S. Manay, "Using Particle Filters to Detect Polygon Structures of Fixed Shape in Still Images", Int. Conf. Signal and Image Proc., August 18-20, 2008, Kailua-Kona, Hawaii.

[8] B. Ristic, S. Arulampalam and N. Gordon, **Beyond the Kalman Filter: Particle Filters for Tracking Applications**, Artech House, 2004, Chapter 3.

[9] Y. Shan and Z. Zhang, "Corner Guided Curve Matching and its Application to Scene Reconstruction", CVPR 2000, vol.2, pp.796-803.

[10] B. Song, I. Yun and S. Lee, "A Target Recognition Technique Employing Geometric Invariants", Pattern Recognition, vol.33, 2000, pp.413-425.